# National Center for Computational Sciences Snapshot
# The Week of September 1, 2008

**Supercomputing Simplified**

*ADIOS boosts big science*

"In computing, things need to be flexible and easy to use." So goes the motto of National Center for Computational Sciences (NCCS) researcher Scott Klasky. He and Georgia Tech's Jay Lofstead and Karsten Schwan are the brains behind a new input/output (I/O) componentization library dubbed the Adapatable I/O System or, more popularly, ADIOS. As supercomputers become faster and more complex, few things are getting simpler. Enter ADIOS.

The middleware's primary goal is to make the process of getting information in and out of a supercomputer easier and more effective. ADIOS's simple application programming interfaces (APIs) and external metadata (XML) file give researchers fast, portable performance, making the choice between efficient I/O and usable data a thing of the past.

Researchers regularly find themselves having to choose between the performance of their applications and the amount and quality of the data they write. It's a problem familiar to Klasky from his early years as a fusion researcher at Princeton Plasma Physics Laboratory using the Gyrokinetic Toroidal Code (GTC), developed by a team led by Zhihong Lin (now at the University of California, Irvine).

"We looked at the performance of how often we would like to write, and we were spending over 20 percent of the time writing the analysis files. Well, 20 percent of all your computational time writing analysis files is too much. The scientists eventually decided that unless it was a run that we definitely wanted to get some visualization out of, we weren't going to write those because we were wasting our valuable computing time doing this." Now researchers using ADIOS can run high-performance applications and write high-quality data, with greatly improved I/O rates.

In a recent fusion simulation on the NCCS's flagship Cray XT4, known as Jaguar, researchers using GTC immediately saw huge I/O benefits in a simulation that produced 60 terabytes of data. "With ADIOS, the I/O was vastly improved, consuming less than 3 percent of run time and allowing the researchers to write tens of terabytes of data smoothly without file system failure," said principal investigator Yong Xiao of the University of California-Irvine. "This huge amount of data needs fast and smooth file writing and reading. With poor I/O, the file writing takes up precious computer time and the parallel file systems on high-performance computing systems can choke."

ADIOS has also made huge strides with Chimera, an astrophysics code. During recent runs on Jaguar the necessary I/O time went from approximately 20 minutes to 1.4 seconds, an improvement of approximately1,000 times Whatever the application, ADIOS's I/O implementation is given at runtime and can be different for different types of output, such as checkpoint restart files and three-dimensional diagnostics, adding to the middleware's I/O arsenal.

Perhaps just as important to researchers as improved I/O is the fact that ADIOS is inserted with an application's XML files, leaving the application's source code untouched. Metadata contains information about all of the variables in a simulation and, among other things, helps the researchers understand what they are looking at when they review data a week, month, or year down the road. Anyone who has ever stared at a year-old file and wondered what it was good for will understand. And when you're dealing with terabytes of data, the problem only gets worse.

Traditionally, programmers would have to encode all of their metadata directly in their source code. Once a code was deemed stable, programmers were reluctant to make any changes that might affect its proper operation. ADIOS's XML file captures all of that information, allowing late additions of rich metadata without the threat of inadvertently introducing bugs. Additionally, ADIOS stores the metadata along with the data in the output, making it easier for scientists to understand exactly what the data represents.

Finally, by supporting asynchronous transport methods, ADIOS allows for even faster I/O by, again, simply switching out an XML file. The trick is scheduling. I/O uses the network bandwidth, and by taking advantage of the downtime during communication between processors, researchers "can essentially get I/O for free," said Klasky.

Another boon to researchers is ADIOS's ability to provide on-the-spot data analysis. For example, it allows researchers to perform real-time (in situ) visualization by inserting the API into their source code and simply altering the XML file. And ADIOS offers the ability to perform data multiplexing, if needed, enabling a researcher to perform both in situ visualization and I/O simultaneously.

In addition to its simplicity and speed, ADIOS is likewise becoming increasingly scalable. Klasky expects it to scale up to 120,000 processors on the Cray XT4 in the future. So far, its largest run was the aforementioned GTC simulation, running on 29,000 processors with 3 percent overhead. And while ADIOS is currently compatible with only Fortran, C, and C++, Klasky recognizes the need for it to operate smoothly on multiple architectures, as well as with different versions of software. His team is working to extend ADIOS to as many systems and applications as possible. To date it has been validated on the NCCS's Jaguar and on the center's Linux clusters. Later in the year Klasky expects to extend ADIOS's reach to IBM Blue Gene supercomputers such as the NCCS's Eugene system.

Eventually, said Klasky, his team wants to release the software as open source. While this means more work for the team—documentation, tutorials, bug searchers, etc.—it will also accelerate ADIOS development.

"We're trying to say, 'Here are our initial codes,'" Klasky explained. "Let's make sure they run on all these different architectures . . . Eventually we'll need people to put it into their stuff instead of us putting it into their codes."

There are few foreseeable limits to ADIOS's potential. As it is expanded to additional platforms, simulating big science will become correspondingly simpler, allowing researchers to concentrate more on their results than the technical aspects of their simulations. And as high-performance computing (HPC) becomes an increasingly powerful research tool, there will be no shortage of grateful scientists.

For more information on ADIOS, visit http://www.nccs.gov/wp-content/training/2008_users_meeting/4-17-08/VV-Petascale-4-17-08.pdf

**Tap It and Trap It**

*Groundwater simulation addresses the challenges of carbon sequestration*

Not every response to global warming focuses on new energy sources. Even as we develop promising technologies such as solar power, biofuels, and nuclear energy, we face the prospect of being tethered for some time to the old energy sources—primarily fossil fuels such as coal and oil.

One proposal for mitigating the effect of coal power on the earth's climate involves separating carbon dioxide—or $CO_2$—from power plant emissions and pumping it deep underground, where it can remain indefinitely dissolved in the groundwater or converted into a solid form of carbonate minerals.

A team of researchers led by Peter Lichtner of Los Alamos National Laboratory (LANL) is using the National Center for Computational Sciences' Jaguar supercomputer to simulate this process, known as carbon sequestration, searching for ways to maximize the benefits and avoid potential drawbacks. Using Jaguar, the team has been able to conduct the largest groundwater simulations ever seen, pursuing its research with a 2008 DOE/Office of Advanced Scientific Computing Research Joule metric code known as PFLOTRAN.

Coal is very abundant in the United States, but coal power carries a variety of serious problems, one of which is that coal-fired power plants spew $CO_2$ into the air. The process being simulated by Lichtner's team involves taking $CO_2$ that has been separated from a power plant's emissions and injecting it nearby into a deep saline aquifer one to two kilometers below the surface. If all goes according to plan, it would spread out under a layer of impermeable rock and get an opportunity to dissolve into the surrounding brine.

The $CO_2$ would be pumped in a state known as a supercritical phase, which is present when it is kept above 50 degrees centigrade—120 degrees Fahrenheit—and over 100 times atmospheric pressure; it would be kept in that state by the heat and pressure naturally present deep underground. According to Lichtner, $CO_2$ in this phase is in some ways like a liquid and in some like a gas, but the primary benefit is that it avoids the rapid expansion that would go along with changes between the two phases.

Lichtner's team is investigating a process known as fingering that speeds the rate at which the $CO_2$ dissolves. Fingering grows out of the fact that while $CO_2$ in the supercritical phase is lighter than the surrounding brine, brine in which $CO_2$ has been dissolved is actually heavier than unsaturated brine. The result is a convection current, with fingers of the heavier, saturated brine sinking. This fingering in turn increases the surface area between the $CO_2$ and the brine and speeds the dissolution of the supercritical $CO_2$ into the brine.

The rate of dissolution is critical to the success of carbon sequestration. When it is first injected in the ground, the $CO_2$ pushes the brine out of place. Once the $CO_2$ dissolves, however, it adds little to the volume of the brine, which can then move back into place.

"The problem is that we're talking about injecting huge amounts of $CO_2$ by volume," Lichtner explained. "If you were injecting it into a deep saline aquifer, for example, you would initially have to displace the brine that was present, and then the question is, 'Where does that go?' It's a race against time how rapidly this $CO_2$ will dissipate."

There are other hazards as well that must be thoroughly understood before large volumes of $CO_2$ can be pumped underground. If the $CO_2$ were to rise to the surface, that would create another substantial hazard. The process of dissolving $CO_2$ into groundwater is, in fact, known as carbonation; $CO_2$ rising rapidly to the surface could literally turn the groundwater into seltzer water.

"As long as the supercritical phase still exists, it presents a hazard to people living on the surface," Lichtner said, "because it could escape through fractures, abandoned boreholes, or boreholes that leak. And so the rate of dissipation is important to understand to know how rapidly you get rid of this supercritical phase."

A final issue that must be studied focuses not so much on the rate that $CO_2$ dissolves, but rather on the changes this process brings to the aquifer itself. As Lichtner explained, $CO_2$ produces carbonic acid, which in turn lowers the pH of the brine. This could speed the reaction between the newly acidic brine and surrounding minerals and potentially release contaminants into the environment that would not be present otherwise.

Lichtner's team is simulating carbon sequestration using an application known as PFLOTRAN, which is built on the PETSc parallel libraries developed by a team led by Barry Smith at Argonne National Laboratory. Chuan Lu of the University of Utah developed the supercritical $CO_2$ implementation in PFLOTRAN while working with Lichtner as a postdoctoral researcher at LANL. Lichtner and his team have shown that PFLOTRAN can handle grids on the order of a billion cells—an unprecedentedly large number for a groundwater simulation. Nevertheless, each cell in such a simulation will be nearly 100 square meters, too large to comfortably analyze the fingering process, which takes place at the scale of tens of centimeters to tens of meters, depending on the properties of the aquifer.

Lichtner noted that his team is working both to improve the performance of PFLOTRAN and to prepare for the arrival of even more powerful supercomputers. To make PFLOTRAN more effective, for example, the team is working to evolve from the use of a structured grid, in which a quarter of the cells give no useful information, to an unstructured grid that can redistribute those cells where they will be of most value.

Beyond that, he said, the team is anticipating a new generation of supercomputers capable of speeds greater than 1,000 trillion calculations a second, or a petaflop. At that scale, Lichtner's team will approach the resources it needs to guide the process of carbon sequestration, enabling us to remove a substantial chunk of $CO_2$ from the atmosphere.

**Smoky Stepping Stone to Big Science**

*System proves its worth in porting, debugging*

The Cray XT4 known as Jaguar may be the NCCS's flagship system, but the big cat gets plenty of big help.

The NCCS's development resource known as Smoky (named after the mascot for the University of Tennessee) was purchased to offload debugging and development work from Jaguar to a smaller yet capable system and is now the primary platform used by researchers to port and scale their codes to Jaguar, the production system. Simply put, if you're developing codes for leadership class computing, you need leadership class computing resources. In its own right, Smoky is very capable of producing big science solo, but its relationship works well with Jaguar primarily because the two programming environments closely mirror one another.

"It's the bullpen for science, getting the players ready for the big game," said NCCS staff member Bobby Whitten. "Without Smoky, development would be seriously slowed down, as everyone would have to wait in line on Jaguar."

Smoky is available to all Innovative and Novel Computational Impact on Theory and Experiment (also known as INCITE) projects, but access is limited to a specific number of people on each team, namely those involved in serious scientific application development. An 11.7 TF system, Smoky consists of 320 quad-core AMD processors and 2.5 terabytes of memory. Smoky also features a gigabit Ethernet network with infiniband interconnect and access to Spider, the NCCS's Lustre-based file system.

If not for Smoky, the big science being conducted on Jaguar would take much longer to accomplish. And when it comes to America's future in areas such as energy, climate, and technology, there is little time to spare.

**Business Liaison Joins ORNL**

*Tichenor has long experience creating HPC partnerships*

Suzy Tichenor, recently of the Washington-based Council on Competitiveness, has joined Oak Ridge National Laboratory's (ORNL's) Computing and Computational Sciences Directorate as director of its Industrial Partnership Program.

Tichenor comes to ORNL from the Washington-based council, where she served as vice president and director of its HPC initiative. She will be the principal interface between the ORNL organization and industry.

“Suzy has a unique understanding of how industry is using high-performance computing to drive innovation and productivity,” said Thomas Zacharia, ORNL associate laboratory director for computing and computational sciences. “She will bring tremendous energy and focus to our industrial outreach activities.”

Tichenor’s appointment underscores the laboratory’s commitment to working with industry, leveraging ORNL’s extensive computing and scientific expertise to solve the most important problems facing business while continuing to advance the frontiers of research and discovery. In her new role, she will lend support to Tom Ballard, director of ORNL’s Partnerships Directorate and responsible for the laboratory’s technology-transfer and economic-development agenda, Zacharia said. Tichenor will work out of UT-Battelle’s Washington, DC, office.

Tichenor has more than 20 years’ experience creating partnerships and programs from within the government, the private sector, and not-for-profit organizations. She has served as the principal investigator for HPC-related grants from the Department of Energy’s Office of Science and National Nuclear Security Administration, the National Science Foundation, and the Department of Defense’s Defense Advanced Research Projects Agency.

She previously held senior positions at Cray Research, Inc., now Cray Inc.; at a startup health-care firm; and at a national not-for-profit organization.